

Büchi Automata can have Smaller Quotients

Lorenzo Clemente

LFCS, School of Informatics, University of Edinburgh, UK

Abstract. We study novel simulation-like preorders for quotienting nondeterministic Büchi automata. We define *fixed-word delayed simulation*, a new preorder coarser than delayed simulation. We argue that fixed-word simulation is the coarsest forward simulation-like preorder which can be used for quotienting Büchi automata, thus improving our understanding of the limits of quotienting. Also, we show that computing fixed-word simulation is PSPACE-complete.

On the practical side, we introduce *proxy simulations*, which are novel polynomial-time computable preorders sound for quotienting. In particular, *delayed proxy simulation* induce quotients that can be smaller by an arbitrarily large factor than direct backward simulation. We derive proxy simulations as the product of a theory of refinement transformers: A *refinement transformer* maps preorders non-decreasingly, preserving certain properties. We study under which general conditions refinement transformers are sound for quotienting.

1 Introduction

Büchi automata minimization is an important topic in automata theory, both for the theoretical understanding of automata over infinite words and for practical applications. Minimizing an automaton means reducing the number of its states as much as possible, while preserving the recognized language. Minimal automata need not be unique, and their structure does not necessarily bear any resemblance to the original model; in the realm of infinite words, this holds even for deterministic models. This hints at why exact minimization has high complexity: Indeed, minimality checking is PSPACE-hard for nondeterministic models (already over finite words [12]), and NP-hard for deterministic Büchi automata [19]. Moreover, even approximating the minimal model is hard [8].

By posing suitable restrictions on the minimization procedure, it is nonetheless possible to trade exact minimality for efficiency. In the approach of *quotienting*, smaller automata are obtained by merging together equivalent states, under appropriately defined equivalences. In particular, quotienting by simulation equivalence has proven to be an effective heuristics for reducing the size of automata in cases of practical relevance.

The notion of *simulation preorder* and *equivalence* [17] is a crucial tool for comparing the behaviour of systems. It is best described via a game between two players, Duplicator and Spoiler, where the former tries to stepwise match the moves of the latter. But not every simulation preorder can be used for quotienting: We call a preorder *good for quotienting* (GFQ) if the quotient automaton (w.r.t. the induced equivalence) recognizes the same language as the original automaton. In particular, a necessary condition for a simulation to be GFQ is to take into account the acceptance condition: For example, in *direct simulation* [4], Duplicator has the additional requirement to visit an accepting state whenever Spoiler does so, while in the coarser *fair simulation* [10], Duplicator has to visit infinitely many accepting states if Spoiler does so. But, while direct

simulation is GFQ [2], fair simulation is not [11].¹ This prompted the development of *delayed* simulation [6], a GFQ preorder intermediate between direct and fair simulation.

We study the border of GFQ preorders. In our first attempt we generalize delayed simulation to *delayed containment*. While in simulation the two players take turns in selecting transitions, in containment the game ends in one round: First Spoiler picks an infinite path, and then Duplicator has to match it with another infinite path. The winning condition is delayed-like: Every accepting state of Spoiler has to be matched by an accepting state of Duplicator, possibly occurring later. Therefore, in delayed containment Duplicator is much stronger than in simulation; in other words, containment is coarser than simulation. In fact, it is *too coarse*: We give a counterexample where delayed containment is not GFQ. We henceforth turn our attention to finer preorders.

In our second attempt, we remedy to the deficiency above by introducing *fixed-word delayed simulation*, an intermediate notion between simulation and containment. In fixed-word simulation, Spoiler does not reveal the whole path in advance like in containment; instead, she only declares the input word beforehand. Then, the simulation game starts, but now transitions can be taken only if they match the word fixed earlier by Spoiler. Unlike containment, fixed-word delayed simulation is GFQ, as we show.

We proceed by looking at even coarser GFQ preorders. We enrich fixed-word simulation by allowing Duplicator to use *multiple pebbles*, in the style of [5]. The question arises as whether Duplicator gains more power by “hedging her bets” when she already knows the input word in advance. By using an ordinal ranking argument (reminiscent of [15]), we establish that this is not the case, and that the multi-pebble hierarchy collapses to the 1-pebble case, i.e., to fixed-word delayed simulation itself. Incidentally, this also shows that the whole delayed multi-pebble hierarchy from [5] is entirely contained in fixed-word delayed simulation—the containment being strict.

For what concerns the complexity of computing fixed-word simulation, we establish that it is PSPACE-complete, by a mutual reduction from Büchi automata universality.

With the aim of getting tractable preorders, we then look at a different way of obtaining GFQ relations, by introducing a theory of refinement transformers: A *refinement transformer* maps a preorder \preceq to a coarser preorder \preceq' , s.t., once \preceq is known, \preceq' can be computed with only a polynomial time overhead. The idea is to play a simulation-like game, where we allow Duplicator to “jump” to \preceq -bigger states, called *proxies*, after Spoiler has selected her transition. Duplicator can then reply with a transition from the proxy instead of the original state. We say that proxy states are *dynamic* in the sense that they depend on the transition selected by Spoiler.² Under certain conditions, we show that refinement transformers induce GFQ preorders.

Finally, we introduce *proxy simulations*, which are novel polynomial time GFQ preorders obtained by applying refinement transformers to a concrete preorder \preceq , namely, to *backward* direct simulation (called reverse simulation in [20]). We define two versions of proxy simulation, direct and delayed, the latter being coarser than the former, and both coarser than direct backward simulation. Moreover, we show that the delayed variant can achieve quotients smaller than direct proxy simulation by an arbitrarily large factor. Full proofs can be found in the appendix.

¹ In fact, for Büchi automata it is well-known that also language equivalence is not GFQ.

² Proxies are strongly related to mediators [1]. We compare them in depth in Section 6.

Related work. Delayed simulation [6] has been extended to generalized automata [13], to multiple pebbles [5], to alternating automata [7] and to the combination of the last two [3]. Fair simulation has been used for state space reduction in [9]. The abstract idea of mixing forward and backward modes in quotienting can be traced back at least to [18]; in the context of alternating automata, it has been studied in [1].

2 Preliminaries

Games. For a finite sequence $\pi = e_0 e_1 \cdots e_{k-1}$, let $|\pi| = k$ be its length, and let $\text{last}(\pi) = e_{k-1}$ be its last element. If π is infinite, then take $|\pi| = \omega$.

A *game* is a tuple $G = (P, P_0, P_1, p_I, \Gamma, \Gamma_0, \Gamma_1, W)$, where P is the set of positions, partitioned into disjoint sets P_0 and P_1 , $p_I \in P_0$ is the initial position, $\Gamma = \Gamma_0 \cup \Gamma_1$ is the set of moves, where $\Gamma_0 \subseteq P_0 \times P_1$ and $\Gamma_1 \subseteq P_1 \times P_0$ are the set of moves of Player 0 and Player 1, respectively, and $W \subseteq P_0^\omega$ is the winning condition. A *path* is a finite or infinite sequence of states $\pi = p_0^0 p_0^1 p_1^1 \cdots$ starting in p_I , such that, for all $i < |\pi|$, $(p_i^0, p_{i+1}^1) \in \Gamma_0$ and $(p_i^1, p_{i+1}^0) \in \Gamma_1$. *Partial plays* and *plays* are finite and infinite paths, respectively. We assume that there are no dead ends in the game. A play is *winning* for Player 1 iff $p_0^0 p_1^0 p_2^0 \cdots \in W$; otherwise, it is winning for Player 0.

A *strategy* for Player 0 is a partial function $\sigma_0 : (P_0 P_1)^* P_0 \mapsto P_1$ s.t., for any partial play $\pi \in (P_0 P_1)^* P_0$, if σ_0 is defined on π , then $\pi \cdot \sigma_0(\pi)$ is again a partial play. A play π is σ_0 -conform iff, for every $i \geq 0$, $p_i^1 = \sigma_0(p_0^0 p_0^1 \cdots p_i^0)$. Similarly, a strategy for Player 1 is a partial function $\sigma_1 : (P_0 P_1)^+ \mapsto P_0$ s.t., for any partial play $\pi \in (P_0 P_1)^+$, if σ_1 is defined on π , then $\pi \cdot \sigma_1(\pi)$ is again a partial play. A play π is σ_1 -conform iff, for every $i \geq 0$, $p_{i+1}^0 = \sigma_1(p_0^0 p_0^1 \cdots p_i^1)$. While we do not require strategies to be total functions, we do require that a strategy σ is defined on all σ -conform partial plays.

A strategy σ_i is a *winning strategy* for Player i iff all σ_i -conform plays are winning for Player i . We say that Player i wins the game G if she has a winning strategy.

Automata. A *nondeterministic Büchi automaton* (NBA) is a tuple $\mathcal{Q} = (Q, \Sigma, I, \Delta, F)$, where Q is a finite set of states, Σ is a finite alphabet, $I \subseteq Q$ is the set of initial states, $F \subseteq Q$ is the set of final states and $\Delta \subseteq Q \times \Sigma \times Q$ is the transition relation. We also write $q \xrightarrow{a} q'$ instead of $(q, a, q') \in \Delta$, and just $q \longrightarrow q'$ when $\exists a \in \Sigma \cdot q \xrightarrow{a} q'$. For two sets of states $\mathbf{q}, \mathbf{q}' \subseteq Q$, we write $\mathbf{q} \xRightarrow{a} \mathbf{q}'$ iff $\forall q' \in \mathbf{q}' \cdot \exists q \in \mathbf{q} \cdot q \xrightarrow{a} q'$.³ For a state $q \in Q$, let $[q \in F] = 1$ if q is accepting, and 0 otherwise. We assume that every state is reachable from some initial state, and that the transition relation is total.

For a finite or infinite sequence of states $\rho = q_0 q_1 \cdots$ and an index $i \leq |\rho|$, let $\text{cnt-final}(\rho, i)$ be the number of final states occurring in ρ up to (and including) the i -th element. Formally, $\text{cnt-final}(\rho, i) = \sum_{0 \leq k \leq i} [q_k \in F]$, with $\text{cnt-final}(\rho, 0) = 0$. Let $\text{cnt-final}(\rho) = \text{cnt-final}(\rho, |\rho|)$. If ρ is infinite, then $\text{cnt-final}(\rho) = \omega$ iff ρ contains infinitely many accepting states.

Fix a finite or infinite word $w = a_0 a_1 \cdots$. A *path* π over w is a sequence $q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \cdots$ of length $|w| + 1$. A path is *initial* if it starts in an initial state $q_0 \in I$, it is a *run* if it is initial and infinite, and it is *fair* if $\text{cnt-final}(\pi) = \omega$. An *accepting run* is a run which is fair. The *language* $\mathcal{L}^\omega(\mathcal{Q})$ of a NBA \mathcal{Q} is the set of infinite words which admit an accepting run, i.e., $\mathcal{L}^\omega(\mathcal{Q}) = \{w \in \Sigma^\omega \mid \text{there exists an accepting run } \pi \text{ over } w\}$.

³ This kind of backward-compatible transition had already appeared in [16].

Quotients. Let $\mathcal{Q} = (Q, \Sigma, I, \Delta, F)$ be a NBA and let R be any binary relation on Q . We say that \approx_R is the *equivalence induced by R* if \approx_R is the largest equivalence contained in the transitive and reflexive closure of R . I.e., $\approx_R = R^* \cap (R^*)^{-1}$. Let the function $[\cdot]_R : Q \mapsto 2^Q$ map each element $q \in Q$ to the equivalence class $[q]_R \subseteq Q$ it belongs to, i.e., $[q]_R := \{q' \in Q \mid q \approx_R q'\}$. We overload $[P]_R$ on sets $P \subseteq Q$ by taking the set of equivalence classes. When clear from the context, we avoid noting the dependence of \approx and $[\cdot]$ on R .

An equivalence \approx on \mathcal{Q} induces the *quotient automaton* $\mathcal{Q}_{\approx} = ([Q], \Sigma, [I], \Delta_{\approx}, [F])$, where, for any $q, q' \in Q$ and $a \in \Sigma$, $([q], a, [q']) \in \Delta_{\approx}$ iff $(q, a, q') \in \Delta$. This is called a *naïve* quotient since both initial/final states and transitions are induced representative-wise. When we quotient w.r.t. a relation R which is not itself an equivalence, we actually mean quotienting w.r.t. the induced equivalence \approx . We say that R is *good for quotienting* (GFQ) if quotienting \mathcal{Q} w.r.t. R preserves the language, that is, $\mathcal{L}^\omega(\mathcal{Q}) = \mathcal{L}^\omega(\mathcal{Q}_{\approx})$.

Lemma 1. *For two equivalences \approx_0, \approx_1 , if $\approx_0 \subseteq \approx_1$, then $\mathcal{L}^\omega(\mathcal{Q}_{\approx_0}) \subseteq \mathcal{L}^\omega(\mathcal{Q}_{\approx_1})$. In particular, by letting \approx_0 be the identity, $\mathcal{L}^\omega(\mathcal{Q}) \subseteq \mathcal{L}^\omega(\mathcal{Q}_{\approx_1})$.*

3 Quotienting with forward simulations

In this section we study several generalizations of delayed simulation, in order to investigate the border of good for quotienting (GFQ) forward-like preorders. In our first attempt we introduce *delayed containment*, which is obtained as a modification of the usual simulation interaction between players: In the delayed containment game between q and s there are only two rounds. Spoiler moves first and selects both an infinite word $w = a_0 a_1 \dots$ and an infinite path $q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} \dots$ over w starting in $q = q_0$; then, Duplicator replies with an infinite path $s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} \dots$ over w starting in $s = s_0$. The winning condition is delayed-like: $\forall i \cdot q_i \in F \implies \exists j \geq i \cdot s_j \in F$. If Duplicator wins the delayed containment game between q and s , we write $q \subseteq^{\text{de}} s$. Clearly, \subseteq^{de} is a preorder implying language containment. One might wonder whether delayed-containment is GFQ. Unfortunately, this is not the case (see Figure 5 in the Appendix). Therefore, \subseteq^{de} is too coarse for quotienting, and we shall look at finer relations.

Lemma 2. *\subseteq^{de} is not a GFQ preorder.*

3.1 Fixed-word delayed simulation

Our second attempt at generalizing delayed simulation still retains the flavour of containment. While in containment \subseteq^{de} Spoiler reveals both the input word w and a path over w , in *fixed-word* simulation $\sqsubseteq_{\text{fx}}^{\text{de}}$ Spoiler reveals w only. Then, after w has been fixed, the game proceeds like in delayed simulation, with the proviso that transitions match symbols in w .⁴ Formally, let $w = a_0 a_1 \dots \in \Sigma^\omega$. In the w -simulation game $G_w^{\text{de}}(q, s)$ the set of positions of Spoiler is $P_0 = Q \times Q \times \mathbb{N}$, the set of positions of Duplicator is $P_1 = Q \times Q \times Q \times \mathbb{N}$ and $\langle q, s, 0 \rangle$ is the initial position. Transitions are determined as follows: Spoiler can select a move of the form $(\langle q, s, i \rangle, \langle q, s, q', i \rangle) \in \Gamma_0^{w\text{-de}}$

⁴ The related notion of fixed-word *fair* simulation clearly coincides with ω -language inclusion.

if $q \xrightarrow{a_i} q'$, and Duplicator can select a move of the form $(\langle q, s, q', i \rangle, \langle q', s', i + 1 \rangle) \in \Gamma_1^{w\text{-de}}$ if $s \xrightarrow{a_i} s'$. Notice that the input symbol a_i is fixed, and it has to match the corresponding symbol in w . The winning condition is $W = \{ \langle q_0, s_0, 0 \rangle \langle q_1, s_1, 1 \rangle \cdots \mid \forall i. q_i \in F \implies \exists j \geq i. s_j \in F \}$. Let $q \sqsubseteq_w^{\text{de}} s$ iff Duplicator wins the w -simulation game $G_w^{\text{de}}(q, s)$, and $q \sqsubseteq_{\text{fx}}^{\text{de}} s$ iff $q \sqsubseteq_w^{\text{de}} s$ for all $w \in \Sigma^\omega$. Clearly, fixed-word simulation is a preorder implying containment.

Fact 1. $\sqsubseteq_{\text{fx}}^{\text{de}}$ is a reflexive and transitive relation, and $\forall q, s \in Q. q \sqsubseteq_{\text{fx}}^{\text{de}} s \implies q \sqsubseteq^{\text{de}} s$.

Unlike delayed containment, fixed-word delayed simulation is GFQ. Moreover, fixed-word delayed simulation quotients can be more succinct than (multi)pebble delayed simulation quotients by an arbitrarily large factor. See Figure 6 in the Appendix.

Theorem 1. $\sqsubseteq_{\text{fx}}^{\text{de}}$ is good for quotienting.

Complexity of delayed fixed word simulation. Let q, s be two states in Q . We reduce the problem of checking $q \sqsubseteq_{\text{fx}}^{\text{de}} s$ to the universality problem of a suitable alternating Büchi product automaton (ABA) \mathcal{A} . We design \mathcal{A} to accept exactly those words w s.t. Duplicator wins $G_w^{\text{de}}(q, s)$. Then, by the definition of $\sqsubseteq_{\text{fx}}^{\text{de}}$, it is enough to check whether \mathcal{A} has universal language. See [21] (or Appendix A.1) for background on ABAs.

The idea is to enrich configurations in the fixed-word simulation game by adding an obligation bit recording whether Duplicator has any pending constraint to visit an accepting state. Initially the bit is 0, and it is set to 1 whenever Spoiler is accepting; a reset to 0 can occur afterwards, if and when Duplicator visits an accepting state.

Let $\mathcal{Q} = (Q, \Sigma, I, \Delta, F)$ be a NBA. We define a product ABA $\mathcal{A} = (A, \Sigma, \delta, \alpha)$ as follows: The set of states is $A = Q \times Q \times \{0, 1\}$, final states are of the form $\alpha = Q \times Q \times \{0\}$ and, for any $\langle q, s, b \rangle \in A$ and $a \in \Sigma$,

$$\delta(\langle q, s, b \rangle, a) = \bigwedge_{q \xrightarrow{a} q'} \bigvee_{s \xrightarrow{a} s'} \langle q', s', b' \rangle, \quad \text{where } b' = \begin{cases} 0 & \text{if } s \in F \\ 1 & \text{if } q \in F \wedge s \notin F \\ b & \text{otherwise} \end{cases}$$

It follows directly from the definitions that $q \sqsubseteq_{\text{fx}}^{\text{de}} s$ iff $\mathcal{L}^\omega(\langle q, s, 0 \rangle) = \Sigma^\omega$. A reduction in the other direction is immediate already for NBAs: In fact, an NBA \mathcal{Q} is universal iff $\mathcal{U} \sqsubseteq_{\text{fx}}^{\text{de}} \mathcal{Q}$, where \mathcal{U} is the trivial, universal one-state automaton with an accepting Σ -loop. It is well-known that universality is PSPACE-complete for ABAs/NBAs [14].

Theorem 2. Computing fixed-word delayed simulation is PSPACE-complete.

3.2 Multi-pebble fixed-word delayed simulation

Having established that fixed-word simulation is GFQ, the next question is whether we can find other natural GFQ preorders between fixed-word and delayed containment. A natural attempt is to add a multi-pebble facility on top of $\sqsubseteq_{\text{fx}}^{\text{de}}$. Intuitively, when Duplicator uses multiple pebbles she can “hedge her bets” by moving pebbles to several successors. This allows Duplicator to delay committing to any particular choice by arbitrarily many steps: In particular, she can always gain knowledge on any *finite* number of moves by Spoiler. Perhaps surprisingly, we show that *Duplicator does not gain more*

power by using pebbles. This is stated in Theorem 3, and it is the major technical result of this section. It follows that, once Duplicator knows the input word in advance, there is no difference between knowing only the next step by Spoiler, or the next l steps, for any finite $l > 1$. Yet, if we allow $l = \omega$ lookahead, then we recover delayed containment \sqsubseteq^{de} , which is not GFQ by Lemma 2. Therefore, w.r.t. to the degree of lookahead, $\sqsubseteq_{\text{fx}}^{\text{de}}$ is the coarsest GFQ relation included in \sqsubseteq^{de} .

We now define the multi-pebble fixed-word delayed simulation. Let $k \geq 1$ and $w = a_0 a_1 \dots \in \Sigma^\omega$. In the k -multi-pebble w -delayed simulation game $G_w^{k\text{-de}}(q, s)$ the set of positions of Spoiler is $Q \times 2^Q \times \mathbb{N}$, the set of positions of Duplicator is $Q \times 2^Q \times Q \times \mathbb{N}$, the initial position is $\langle q, \{s\}, 0 \rangle$, and transitions are: $(\langle q, s, i \rangle, \langle q, s, q', i \rangle) \in \Gamma_0$ iff $q \xrightarrow{a_i} q'$, and $(\langle q, s, q', i \rangle, \langle q', s', i+1 \rangle) \in \Gamma_1$ iff $s \xrightarrow{a_i} s'$ and $|s'| \leq k$.

Before defining the winning set we need some preparation. Given an infinite sequence $\pi = \langle q_0, s_0, 0 \rangle \langle q_1, s_1, 1 \rangle \dots$ over $w = a_0 a_1 \dots$ and a round $j \geq 0$, we say that a state $s \in s_j$ *has been accepting* since some previous round $i \leq j$, written $\text{accepting}_j^i(s, \pi)$, iff either $s \in F$, or $i < j$ and there exists $\hat{s} \in s_{j-1}$ s.t. $\hat{s} \xrightarrow{a_{j-1}} s$ and $\text{accepting}_{j-1}^i(\hat{s}, \pi)$. We say that s_j is *good since round* $i \leq j$, written $\text{good}_j^i(s_j, \pi)$, iff at round j every state $s \in s_j$ has been accepting since round i , and j is the least round for which this holds [5]. Duplicator wins a play if, whenever $q_i \in F$ there exists $j \geq i$ s.t. $\text{good}_j^i(s_j, \pi)$. We write $q \sqsubseteq_w^{k\text{-de}} s$ iff Duplicator wins $G_w^{k\text{-de}}(q, s)$, and we write $q \sqsubseteq_{\text{fx}}^{k\text{-de}} s$ iff $\forall w \in \Sigma^\omega. q \sqsubseteq_w^{k\text{-de}} s$.

Clearly, pebble simulations induce a non-decreasing hierarchy: $\sqsubseteq_{\text{fx}}^{1\text{-de}} \subseteq \sqsubseteq_{\text{fx}}^{2\text{-de}} \subseteq \dots$. We establish that the hierarchy actually collapses to the $k = 1$ level. This result is non-trivial, since the delayed winning condition requires reasoning not only about the *possibility* of Duplicator to visit accepting states in the future, but also about exactly *when* such a visit occurs. Technically, our argument uses a ranking argument similar to [15] (see Appendix A.2), with the notable difference that our ranks are *ordinals* ($\leq \omega^2$), instead of natural numbers. We need ordinals to represent how long a player can delay visiting accepting states, and how this events nest with each other. Finally, notice that the result above implies that the multi-pebble delayed simulation hierarchy of [5] is entirely contained in $\sqsubseteq_{\text{fx}}^{\text{de}}$, and the containment is strict (Fig. 6 in the appendix).

Theorem 3. *For any NBA \mathcal{Q} , $k \geq 1$ and states $q, s \in Q$, $q \sqsubseteq_{\text{fx}}^{k\text{-de}} s$ iff $q \sqsubseteq_{\text{fx}}^{\text{de}} s$.*

4 Jumping-safe relations

In this section we present the general technique which is used throughout the paper to establish that preorders are GFQ. We introduce *jumping-safe relations*, which are shown to be GFQ (Theorem 4). In Section 5 we use jumping-safety as an invariant when applying refinement transformers. We start off with an analysis of accepting runs.

Coherent sequences of paths. Fix an infinite word $w \in \Sigma^\omega$. Let $\Pi := \pi_0, \pi_1, \dots$ be an infinite sequence of longer and longer *finite* initial paths in \mathcal{Q} over (prefixes of) w . We are interested in finding a sufficient condition for the existence of an accepting run over w . A necessary condition is that the number of final states in π_i grows unboundedly as i goes to ω . In the case of deterministic automata this condition is also sufficient: Indeed,

in a deterministic automaton there exists a unique run over w , which is accepting exactly when the number of accepting states visited by its prefixes goes to infinity. In this case, we say that the π_i 's are *strongly coherent* since they next path extends the previous one. Unfortunately, in the general case of nondeterministic automata it is quite possible to have paths that visit arbitrarily many final states but no accepting run exists. This occurs because final states can appear arbitrarily late. Indeed, consider Figure 1. Take $w = aba^2ba^3b \dots$: For every prefix $w_i = aba^2b \dots a^i$ there exists a path $\pi_i = qq \dots q \cdot s^i$ over w_i visiting a final state i times. Still, $w \notin \mathcal{L}^\omega(\mathcal{Q})$.




Fig. 1. Automaton \mathcal{Q} .

Therefore, we forbid accepting states to “clump away” in the tail of the path. We ensure this by imposing the existence of an infinite sequence of indices j_0, j_1, \dots s.t., for all i , and for all k_i big enough, the number of final states in π_{k_i} up to the j_i -th state is at least i . In this way, we are guaranteed that at least i final states are present within j_i steps in all but finitely many paths.

Definition 1. Let $\Pi := \pi_0, \pi_1, \dots$ be an infinite sequence of finite paths. We say that Π is a *coherent sequence of paths* if the following property holds:

$$\forall i \cdot \exists j \cdot \exists h \cdot \forall k \geq h \cdot j < |\pi_k| \wedge \text{cnt-final}(\pi_k, j) \geq i. \quad (1)$$

Lemma 3. If Π is coherent, then any infinite subsequence Π' thereof is coherent.

We sketch below the proof that coherent sequences induce fair paths. Let $\Pi = \pi_0, \pi_1, \dots$ be a coherent sequence of paths in \mathcal{Q} . Let $i = 1$, and let j_1 be the index witnessing Π is coherent. Since the π_k 's are branches in a finitely branching tree, there are only a finite number of different prefixes of length j_1 . Therefore, there exists a prefix ρ_1 which is common to infinitely many paths. Let $\Pi' = \pi'_0, \pi'_1, \dots$ be the infinite subsequence of Π containing only suffixes of ρ_1 . Clearly ρ_1 contains at least 1 final state, and each π' in Π' extends ρ_1 . By Lemma 3, Π' is coherent. For $i = 2$, we can apply the reasoning again to Π' , and we obtain a longer prefix ρ_2 extending ρ_1 , and containing at least 2 final states. Let Π'' be the coherent subsequence of Π' containing only suffixes of ρ_2 . In this fashion, we obtain an infinite sequence of *strongly coherent* (finite) paths ρ_1, ρ_2, \dots s.t. ρ_i extends ρ_{i-1} and contains at least i final states. The infinite path to which the sequence converges is the fair path we are after.

Lemma 4. Let $w \in \Sigma^\omega$ and π_0, π_1, \dots as above. If π_0, π_1, \dots is coherent, then there exists a *fair path* ρ over w . Moreover, if all π_i 's are initial, then ρ is initial.

Jumping-safe relations. We established that coherent sequences induce accepting paths. Next, we introduce *jumping-safe* relations, which are designed to induce coherent sequences (and thus accepting paths) when used in quotienting. The idea is to view a path in the quotient automaton as a jumping path in the original automaton, where a “jumping path” is one that can take arbitrary jumps to equivalent states. Jumping-safe relations allows us to transform the sequence of prefixes of an accepting jumping path into a coherent sequence of non-jumping paths; by Lemma 4, this induces a (nonjumping) accepting path.

Fix a word $w = a_0 a_1 \dots \in \Sigma^\omega$, and let R be a binary relation over Q . An R -jumping path is an infinite sequence

$$\pi = q_0 R q_0^F R \hat{q}_0 \xrightarrow{a_0} q_1 R q_1^F R \hat{q}_1 \xrightarrow{a_1} q_2 \dots, \quad (2)$$

and we say that π is *initial* if $q_0 \in I$, and *fair* if $q_i^F \in F$ for infinitely many i 's.

Definition 2. A binary relation R is jumping-safe iff for any initial R -jumping path π there exists an infinite sequence of initial finite paths π_0, π_1, \dots over suitable prefixes of w s.t. $\text{last}(\pi_i) R q_i$ and, if π is fair, then π_0, π_1, \dots is coherent.

Theorem 4. Jumping-safe preorders are good for quotienting.

In Section 5 we introduce refinement transformers, which are designed to preserve jumping-safety. Then, in Section 6 we specialize the approach to *backward direct simulation* $\sqsubseteq_{\text{bw}}^{\text{di}}$ [20], which provides an initial jumping-safe preorder, and which we introduce next: $\sqsubseteq_{\text{bw}}^{\text{di}}$ is the coarsest preorder s.t. $q \sqsubseteq_{\text{bw}}^{\text{di}} s$ implies 1) $\forall (q' \xrightarrow{a} q) \cdot \exists (s' \xrightarrow{a} s) \cdot q' \sqsubseteq_{\text{bw}}^{\text{di}} s'$, 2) $q \in F \implies s \in F$, and 3) $q \in I \implies s \in I$.

Fact 2. $\sqsubseteq_{\text{bw}}^{\text{di}}$ is jumping-safe and computable in polynomial time.

5 Refinement transformers

We study how to obtain GFQ preorders coarser than forward/backward simulation. As a preliminary example, notice that it is not possible to generalize simultaneously both forward and backward simulations. See the counterexample in Fig. 2, where any relation coarser than both forward and backward simulation is not GFQ. Let $\approx_{\text{bw}}^{\text{di}}$ and $\approx_{\text{fw}}^{\text{di}}$ be backward and forward direct simulation equivalence, respectively. We have $q_1 \approx_{\text{bw}}^{\text{di}} q_2 \approx_{\text{fw}}^{\text{di}} q_3$, but “glueing together” q_1, q_2, q_3 would introduce the extraneous word ba^ω . Therefore, one needs to choose whether to extend either forward or backward simulation. The former approach has been pursued in the *mediated preorders* of [1] (in the more general context of alternating automata). Here, we extend backward refinements.




Fig. 2.

We define a *refinement transformer* τ_0 mapping a relation R to a new, coarser relation $\tau_0(R)$. We present τ_0 via a forward direct simulation-like game where Duplicator is allowed to “jump” to R -bigger states—called *proxies*. Formally, in the $\tau_0(R)$ simulation game Spoiler’s positions are in $Q \times Q$, Duplicator’s position are in $Q \times Q \times \Sigma \times Q$ and transitions are as follows: Spoiler picks a transition $(\langle s, q \rangle, \langle s', q', a \rangle) \in \Gamma_0$ simply when $q \xrightarrow{a} q'$, and Duplicator picks a transition $(\langle s, q, a, q' \rangle, \langle s', q' \rangle) \in \Gamma_1$ iff there exists a proxy \hat{s} s.t. $s R \hat{s}$ and $\hat{s} \xrightarrow{a} s'$. The winning condition is: $\forall i \geq 0 \cdot q_i \in F \implies \hat{s}_i \in F$. If Duplicator wins starting from the initial position $\langle s, q \rangle$, we write $s \tau_0(R) q$. (Notice that we swapped the usual order between q and s here.)

Lemma 5. For a preorder R , $R \subseteq R \circ \tau_0(R) \subseteq \tau_0(R)$.

Unfortunately, $\tau_0(R)$ is not necessarily a transitive relation. Therefore, it is not immediately clear how to define a suitable equivalence for quotienting. Figure 2 shows that taking the transitive closure of $\tau_0(R)$ is incorrect—already when R is direct backward simulation $\sqsubseteq_{\text{bw}}^{\text{di}}$: Let $\preceq = \tau_0(\sqsubseteq_{\text{bw}}^{\text{di}})$ and let $\approx = \preceq \cap \preceq^{-1}$. We have $q_3 \approx q_2 \approx q_1 \preceq q_3$, but $q_3 \not\preceq q_1$, and forcing $q_1 \approx q_3$ is incorrect, as noted earlier.

Thus, $\tau_0(R)$ is not GFQ and we need to look at its transitive fragments. Let $T \subseteq \tau_0(R)$. We say that R is *F-respecting* if $q R s \wedge q \in F \implies s \in F$, that T is *self-respecting* if Duplicator wins by never leaving T , that T is *appealing* if transitive and self-respecting, and that T *improves on* R if $R \subseteq T$.

Theorem 5. *Let R a F -respecting preorder, and let $T \subseteq \tau_0(R)$ be an appealing, improving fragment of $\tau_0(R)$. If R is jumping-safe, then T is jumping-safe.*

In particular, by Theorem 4, T is GFQ. Notice that requiring that R is GFQ is not sufficient here, and we need the stronger invariant given by jumping-safety.

Given an appealing fragment $T \subseteq \tau_0(R)$, a natural question is whether $\tau_0(T)$ improves on $\tau_0(R)$, so that τ_0 can be applied repeatedly to get bigger and bigger preorders. We see in the next lemma that this is not the case.

Lemma 6. *For any reflexive R , let $T \subseteq \tau_0(R)$ be any appealing fragment of $\tau_0(R)$. Then, $\tau_0(T) \subseteq \tau_0(R)$.*

Efficient appealing fragments. By Theorems 4 and 5, appealing fragments of τ_0 are GFQ. Yet, we have not specified any method for obtaining these. Ideally, one looks for fragments having maximal cardinality (which yields maximal reduction under quotienting), but finding them is computationally expensive. Instead, we define a new transformer τ_1 which is guaranteed to produce only appealing fragments,⁵ which, while not maximal in general, are maximal amongst all *improving* fragments (Lemma 7).

The reason why $\tau_0(R)$ is not transitive is that only Duplicator is allowed to make “ R -jumps”. This asymmetry is an obstacle to compose simulation games. We recover transitivity by allowing Spoiler to jump as well, thus restoring the symmetry. Formally, the $\tau_1(R)$ simulation game is identical to the one for $\tau_0(R)$, the only difference being that also Spoiler is now allowed to “jump”, i.e., she can pick a transition $(\langle s, q \rangle, \langle s, q, a, q' \rangle) \in \Gamma_0$ iff there exists \hat{q} s.t. $q R \hat{q}$ and $\hat{q} \xrightarrow{a} q'$. The winning condition is: $\forall i \geq 0 \cdot \hat{q}_i \in F \implies \hat{s}_i \in F$. Let $s \tau_1(R) q$ if Duplicator wins from position $\langle s, q \rangle$. It is immediate to see that $\tau_1(R)$ is an appealing fragment of $\tau_0(R)$, and that τ_1 is improving on transitive relations R ’s. Thus, for a preorder R , $R \subseteq \tau_1(R) \subseteq \tau_0(R)$. By Theorems 4 and 5, $\tau_1(R)$ is GFQ (if R is F -respecting).

It turns out that $\tau_1(R)$ is actually the *maximal* appealing, improving fragment of $\tau_0(R)$. This is non-obvious, since the class of appealing T ’s is not closed under union—still, it admits a maximal element. Therefore, τ_1 is an optimal solution to the problem of finding appealing, improving fragments of $\tau_0(R)$.

Lemma 7. *For any R , let $T \subseteq \tau_0(R)$ be any appealing fragment of $\tau_0(R)$. If $R \subseteq T$ (i.e., R is improving), then $T \subseteq \tau_1(R)$.*

⁵ τ_1 needs not be the only solution to this problem: Other ways of obtaining appealing fragments of τ_0 might exist. For this reason, we have given a separate treatment of τ_0 in its generality, together with the general correctness statement (Theorem 5).

5.1 Delayed-like refinement transformers

We show that the refinement transformer approach can yield relations even coarser than τ_1 . Our first attempt is to generalize the direct-like winning condition of τ_0 to a delayed one. Let τ_0^{de} be the same as τ_0 except for the different winning condition, which now is: $\forall i \geq 0 \cdot q_i \in F \implies \exists j \geq i \cdot \hat{s}_j \in F$. Clearly, τ_0^{de} inherits the same transitivity issues of τ_0 . Unfortunately, the approach of taking appealing fragments is not sound here, due to the weaker winning condition. See Figure 7 in the Appendix for a counterexample.

We overcome these issues by dropping τ_0^{de} altogether, and directly generalize τ_1 (instead of τ_0) to a delayed-like notion. The *delayed refinement transformer* τ_1^{de} is like τ_1 , except for the new winning condition: $\forall i \geq 0 \cdot \hat{q}_i \in F \implies \exists j \geq i \cdot \hat{s}_j \in F$. Notice that $\tau_1^{\text{de}}(R)$ is at least as coarse as $\tau_1(R)$, and incomparable with $\tau_0(R)$. Once R is given, $\tau_1^{\text{de}}(R)$ can be computed in polynomial time. See Appendix D.

Lemma 8. *For any R , $\tau_1^{\text{de}}(R)$ is transitive.*

Theorem 6. *If R is a jumping-safe F -respecting preorder, then $\tau_1^{\text{de}}(R)$ is jumping-safe.*

6 Proxy simulations

We apply the theory of transformers from Section 5 to a specific F -respecting preorder, namely backward direct simulation, obtaining *proxy simulations*. Notice that proxy simulation-equivalent states need not have the same language; yet, proxy simulations are GFQ (and computable in polynomial time).

6.1 Direct proxy simulation

Let *direct proxy simulation*, written $\sqsubseteq_{xy}^{\text{di}}$, be defined as $\sqsubseteq_{xy}^{\text{di}} := [\tau_1(\sqsubseteq_{bw}^{\text{di}})]^{-1}$.

Theorem 7. $\sqsubseteq_{xy}^{\text{di}}$ is a polynomial time GFQ preorder at least as coarse as $(\sqsubseteq_{bw}^{\text{di}})^{-1}$.

Proxies vs mediators. Direct proxy simulation and mediated preorder [1] are in general incomparable. While proxy simulation is at least as coarse as backward direct simulation, mediated preorder is at least as coarse as *forward* direct simulation. (We have seen in Section 5 that this is somehow unavoidable, since one cannot hope to generalize simultaneously both forward and backward simulation.)

One notable difference between the two notions is that proxies are “dynamic”, while mediators are “static”: While Dupicator chooses the proxy only *after* Spoiler has selected her move, mediators are chosen uniformly w.r.t. Spoiler’s move.

In Figure 3(a) we show a simple example where $\sqsubseteq_{xy}^{\text{di}}$ achieves greater reduction. Recall that mediated preorder M is always a subset of $\sqsubseteq_{fw}^{\text{di}} \circ (\sqsubseteq_{bw}^{\text{di}})^{-1}$ [1]. In the example, static mediators are just the trivial ones already present in forward simulation. Thus, $\sqsubseteq_{fw}^{\text{di}} \circ (\sqsubseteq_{bw}^{\text{di}})^{-1} = \sqsubseteq_{fw}^{\text{di}}$ and mediated preorder M collapses to forward simulation. On the other side, $p \approx_{xy}^{\text{di}} q$ and $p' \approx_{xy}^{\text{di}} q'_b$. Letting $s = [p, q]$ and $s' = [p', q'_b]$, we obtain the quotient in Figure 3(b).

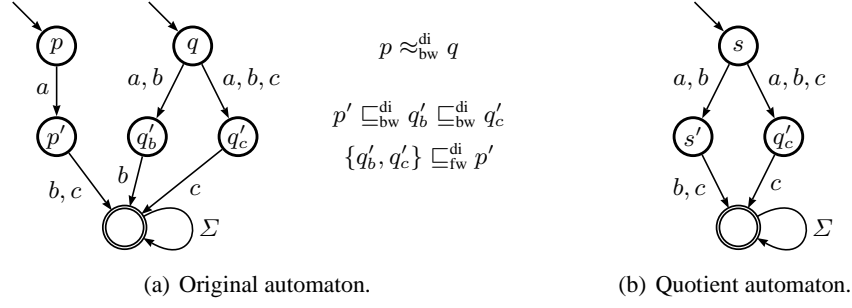


Fig. 3. Direct proxy simulation quotients.

6.2 Delayed proxy simulation

Another difference between the mediated preorder approach [1] and the approach through proxies is that proxies directly enable a delayed simulation-like generalization (see Section 5.1). Again, we fix backward delayed simulation \subseteq_{bw}^{di} as a starting refinement, and we define *delayed proxy simulation* as $\subseteq_{xy}^{de} := [\tau_1^{de}(\subseteq_{bw}^{di})]^{-1}$.

Theorem 8. \subseteq_{xy}^{de} is a polynomial time GFQ preorder.

Notice that delayed proxy simulation is at least as coarse as direct proxy simulation. Moreover, quotients w.r.t. \subseteq_{xy}^{de} can be smaller than direct forward/backward/proxy and delayed simulation quotients by an arbitrary large factor. See Figure 4: Forward delayed simulation is just the identity, and no two states are direct backward or proxy simulation equivalent. But $q_i \subseteq_{bw}^{di} s$ for any $0 < i \leq k-1$. This causes any two outer states q_i, q_j to be \subseteq_{xy}^{de} -equivalent. Therefore, the \subseteq_{xy}^{de} -quotient automaton has only 2 states.




Fig. 4.

7 Conclusions and Future Work

We have proposed novel refinements for quotienting Büchi automata: fixed-word delayed simulation and direct/delayed proxy simulation. Each one has been shown to induce quotients smaller than previously known notions.

We outline a few directions for future work. First, we would like to study practical algorithms for computing fixed-word delayed simulation, and to devise efficient fragments thereof—one promising direction is to look at self-respecting fragments, which usually have lower complexity. Second, we would like to exploit the general correctness argument developed in Section 4 in order to get efficient purely backward refinements (coarser than backward direct simulation). Finally, experiments on cases of practical interest are needed for an empirical evaluation of the proposed techniques.

Acknowledgment. We thank Richard Mayr and Patrick Totzke for helpful discussions, and two anonymous reviewers for their valuable feedback.

References

1. Abdulla, P., Chen, Y.F., Holik, L., Vojnar, T.: Mediating for Reduction. In: FSTTCS. pp. 1–12. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2009)
2. Aziz, A., Singhal, V., Swamy, G.M., Brayton, R.K.: Minimizing Interacting Finite State Machines. Tech. Rep. UCB/ERL M93/68, UoC, Berkeley (1993)
3. Clemente, L., Mayr, R.: Multi- Pebble Simulations for Alternating Automata - (Extended Abstract). In: CONCUR. LNCS, vol. 6269, pp. 297–312. Springer-Verlag (2010), http://dx.doi.org/10.1007/978-3-642-15375-4_21
4. Dill, D.L., Hu, A.J., Wont-Toi, H.: Checking for Language Inclusion Using Simulation Preorders. In: CAV. LNCS, vol. 575. Springer-Verlag (1991), http://dx.doi.org/10.1007/3-540-55179-4_25
5. Etessami, K.: A Hierarchy of Polynomial-Time Computable Simulations for Automata. In: CONCUR. LNCS, vol. 2421, pp. 131–144. Springer-Verlag (2002), http://dx.doi.org/10.1007/3-540-45694-5_10
6. Etessami, K., Wilke, T., Schuller, R.A.: Fair Simulation Relations, Parity Games, and State Space Reduction for Büchi Automata. *SIAM J. Comput.* 34(5), 1159–1175 (2005), <http://epubs.siam.org/sam-bin/dbq/article/42067>
7. Fritz, C., Wilke, T.: Simulation Relations for Alternating Büchi Automata. *Theor. Comput. Sci.* 338(1-3), 275–314 (2005), <http://dx.doi.org/10.1016/j.tcs.2005.01.016>
8. Gramlich, G., Schnitger, G.: Minimizing NFA's and Regular Expressions. *Journal of Computer and System Sciences* 73(6), 908–923 (2007), <http://www.sciencedirect.com/science/article/B6WJ0-4MMP64W-1/2/769c421c0456a3f502681a>
9. Gurumurthy, S., Bloem, R., Somenzi, F.: Fair Simulation Minimization. In: CAV. LNCS, vol. 2404, pp. 610–624. Springer-Verlag (2002), <http://portal.acm.org/citation.cfm?id=647771.734423>
10. Henzinger, T.A., Kupferman, O., Rajamani, S.K.: Fair Simulation. *Information and Computation* 173, 64–81 (2002), <http://dx.doi.org/10.1006/inco.2001.3085>
11. Henzinger, T.A., Rajamani, S.K.: Fair Bisimulation. In: TACAS. LNCS, vol. 1785, pp. 299–314. Springer-Verlag (2000), <http://portal.acm.org/citation.cfm?id=646484.691752>
12. Jiang, T., Ravikumar, B.: Minimal NFA Problems are Hard. In: Albert, J., Monien, B., Artalejo, M. (eds.) *Automata, Languages and Programming, Lecture Notes in Computer Science*, vol. 510, pp. 629–640. Springer Berlin / Heidelberg (1991), http://dx.doi.org/10.1007/3-540-54233-7_169
13. Juvekar, S., Piterman, N.: Minimizing Generalized Büchi Automata. In: CAV. LNCS, vol. 4414, pp. 45–58. Springer-Verlag (2006), http://dx.doi.org/10.1007/11817963_7
14. Kupferman, O., Vardi, M.: Verification of Fair Transition Systems. In: CAV, LNCS, vol. 1102, pp. 372–382. Springer-Verlag (1996), <http://citeseer.ist.psu.edu/viewdoc/summary?doi=10.1.1.29.9654>
15. Kupferman, O., Vardi, M.: Weak Alternating Automata Are Not That Weak. *ACM Trans. Comput. Logic* 2, 408–429 (Jul 2001), <http://dx.doi.org/10.1145/377978.377993>
16. Lynch, N.A., Vaandrager, F.W.: Forward and Backward Simulations. Part I: Untimed Systems. *Information and Computation* 121(2), 214–233 (1995), <http://citeseer.ist.psu.edu/viewdoc/summary?doi=10.1.1.12.3241>
17. Milner, R.: *Communication and Concurrency*. Prentice-Hall (1989)

18. Raimi, R.S.: Environment Modeling and Efficient State Reachability Checking. Ph.D. thesis, The University of Texas at Austin (1999)
19. Schewe, S.: Beyond Hyper-Minimisation—Minimising DBAs and DPAs is NP-Complete. In: Lodaya, K., Mahajan, M. (eds.) FSTTCS. LIPIcs, vol. 8, pp. 400–411. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2010), <http://drops.dagstuhl.de/opus/volltexte/2010/2881>
20. Somenzi, F., Bloem, R.: Efficient Büchi Automata from LTL Formulae. In: CAV, LNCS, vol. 1855, pp. 248–263. Springer-Verlag (2000), http://dx.doi.org/10.1007/10722167_21
21. Vardi, M.: Alternating Automata and Program Verification. In: Computer Science Today, LNCS, vol. 1000, pp. 471–485. Springer-Verlag (1995), <http://dx.doi.org/10.1007/BFb0015261>

A Proofs and additional material for Section 3




Fig. 5. An example showing that delayed containment cannot be employed for quotienting. We have that p_0 is delayed containment equivalent to p_1 . Notice that the automaton \mathcal{A} in (a) does not accept a^ω , but the quotient automaton \mathcal{A}_{\approx} in (b), obtained by identifying p_0 and p_1 , does.

We postpone the proof of Theorem 1 until Section E.




Fig. 6. Fixed-word delayed simulation quotients can achieve arbitrarily high compression ratios.

A.1 Alternating Büchi automata

Below, we give a self-contained definition of alternating Büchi automata. The syntax follows the presentation of [21], while the semantics adheres to [7].

For a set A , let $\mathcal{B}^+(A)$ be the set of positive boolean formulas over A , that is, $\mathcal{B}^+(A)$ is the smallest set containing $A \cup \{\text{true}, \text{false}\}$ and closed under the operations \wedge and \vee . For a formula $\varphi \in \mathcal{B}^+(A)$ and a set $X \subseteq A$, we write $X \models \varphi$ iff the truth assignment assigning **true** to elements in X and **false** to the elements in $A \setminus X$ satisfies φ . An alternating Büchi automaton (ABA) is a tuple $\mathcal{A} = (A, \Sigma, \delta, \alpha)$, where A is a finite set of states, Σ is a finite set of input symbols, $\delta : A \times \Sigma \mapsto \mathcal{B}^+(A)$ is the transition

relation and $\alpha \subseteq A$ is the set of accepting states. Acceptance of an ABA \mathcal{A} is best defined via games [7]. In this context, the two players are usually named Automaton and Pathfinder. Given an infinite word $w = a_0a_1 \dots \in \Sigma^\omega$ and a distinguished starting state p_I , the *acceptance game for w from p_I* is a game where $P_0 = Q \times \omega$ is the set of Automaton's positions, $P_1 = Q \times 2^Q \times \omega$ is the set of Pathfinder's positions, $(p_I, 0)$ is the initial position, and transitions are determined as follows. Automaton can select a transition $(\langle p, i \rangle, \langle p, \mathbf{p}', i \rangle)$ iff $\mathbf{p}' \models \delta(p, a_i)$, and Pathfinder can select a transition $(\langle p, \mathbf{p}', i \rangle, \langle p', i+1 \rangle)$ iff $p' \in \mathbf{p}'$. Finally, the winning condition consists of those paths visiting α infinitely often. A state $p \in A$ accepts $w \in \Sigma^\omega$ iff Automaton wins the acceptance game for w from p . A state p is universal iff it accepts every word $w \in \Sigma^\omega$.

A.2 Proof of Theorem 3

Preliminaries on ordinals. Let ω be the least infinite ordinal, and let ω_1 be the set of all countable ordinals. We denote arbitrary ordinals by α or β , and limit ordinals by λ or μ . In this paper, 0 is considered to be a limit ordinal.

Preliminaries on trees. Let $[n] = \{0, 1, \dots, n-1\}$. A tree domain is a non-empty, prefix-closed subset V of $[n]^*$. With $<_{\text{prf}}$ we denote the prefix order on words; if $u <_{\text{prf}} u'$, then u' is called a descendant of u and u is an ancestor of u' . In particular, if $u' = uc$ for some $c \in \mathbb{N}$, then u' is a child of u . A (labelled) L -tree is a pair (V, t) , where V is a tree domain and $t : V \mapsto L$ is a mapping which assigns a label from L to any node in the tree.

The ranking construction. Let $\mathcal{Q} = (Q, \Sigma, I, \Delta, F)$ be an automaton, and let n be the cardinality of Q . Given an infinite word $w = a_0a_1 \dots \in \Sigma^\omega$, we associate to any state $q \in Q$ a tree domain T_q^w and a Q -tree (T_q^w, t_q^w) , the *unravelling* of Q from q while reading w , by applying the following two rules:

- $\varepsilon \in T_q^w$ and $t_q^w(\varepsilon) = q$.
- If u has length i , $u \in T_q^w$, $t_q^w(u) = p$ and $\Delta(p, a_i) = \{p'_0, p'_1, \dots, p'_{k-1}\}$, then, for any j s.t. $0 \leq j < k$, $uj \in T_q^w$ and $t_q^w(uj) = p'_j$.

It is easy to see that if two nodes at the same level have the same label, then they generate isomorphic subtrees. Therefore, we can “compress” (T_q^w, t_q^w) into an infinite DAG $G_q^w = (V, E)$, where $V \subseteq Q \times \mathbb{N}$ is such that $\langle q, l \rangle \in V$ iff there exists a node in (T_q^w, t_q^w) at level l with label q , and $(\langle q, l \rangle, \langle q', l+1 \rangle) \in E$ iff there exist two nodes u and u' , labelled with q and q' , respectively, s.t. u' is a child of u in (T_q^w, t_q^w) . We say that a vertex $\langle q, l \rangle$ is *accepting* iff $q \in F$.

For any $G \subseteq G_q^w$, we say that a vertex $\langle q, l \rangle$ is a *dead end* in G iff it has no successor in G , and we say that it is *inert* in G iff no accepting vertex can be reached from $\langle q, l \rangle$ in G . In particular, an inert vertex is not accepting. The *girth* of G at level l is the maximal number of vertices of the form $\langle q, l \rangle$ in G , and the *width* of G is the maximal girth over infinitely many levels.

We build a nonincreasing transfinite sequence of DAGs $\{G_\alpha \mid \alpha < \omega_1\}$ as follows:

$$\begin{aligned} G_0 &= G_q^w \\ G_{\alpha+1} &= G_\alpha \setminus \{\langle q, l \rangle \mid \langle q, l \rangle \text{ is a dead end in } G_\alpha\} \\ G_\lambda &= H_\lambda \setminus \{\langle q, l \rangle \mid \langle q, l \rangle \text{ is inert in } H_\lambda\}, \end{aligned}$$

where, for any ordinal α , $H_\alpha = \bigcap_{\beta < \alpha} G_\beta$. Notice that $H_{\alpha+1} = G_\alpha$; and $\alpha \leq \beta$ implies $G_\beta \subseteq G_\alpha$.

Assume that there is no path in G_q^w with an infinite number of accepting vertices. As a direct consequence of König's Lemma, we have that when moving from H_λ to G_λ an infinite path is removed from the graph. Therefore, the width of G_λ is strictly less than the width of H_λ . Since the width of G_q^w is (uniformly) bounded by ω , it follows that H_{ω^2} is empty, and thus G_{ω^2} is empty as well. Therefore, each vertex is either a dead end in G_α or inert in H_λ . In the former case $\langle q, l \rangle$ is in G_α but not in $G_{\alpha+1}$, whereas in the latter case $\langle q, l \rangle$ is in H_λ but not in G_λ . Accordingly, we associate an ordinal *rank* to every vertex $\langle q, l \rangle$ in G_q^w :

$$\text{rank}_q^w(q, l) = \sup_{\alpha < \omega^2} \{ \alpha \mid \langle q, l \rangle \in H_\alpha \}. \quad (\text{Rank})$$

Therefore, under the assumption that G_q^w does not contain any fair path, no vertex receives rank ω^2 . On the other side, if G_q^w contained a fair path, then there exists an infinite path of non-inert vertices starting at $\langle q, 0 \rangle$: In this case, the ranking construction “does not terminate” and stabilizes (at most) at a nonempty $G_{\omega^2} = G_\alpha \neq \emptyset$ for all $\alpha \geq \omega^2$. Thus, vertices in G_{ω^2} would receive rank ω^2 according to (Rank). Since no conflict can arise, we drop any assumption about fair paths thereafter, and we uniformly apply (Rank) in either case.

Remark 1. It is clear from (Rank) that no ordinal larger than ω^2 is actually used in our construction. In fact, we could have given an equivalent presentation in terms of pairs of natural numbers ordered lexicographically. However, we have chosen to use ordinals $\leq \omega^2$ for technical convenience.

Remark 2. A vertex $\langle q, l \rangle$ is in H_α iff it has rank $\geq \alpha$, and it is not in G_α iff it has rank $\leq \alpha$. Therefore, $\text{rank}_q^w(q, l) = \alpha \iff \langle q, l \rangle \in H_\alpha \setminus G_\alpha$.

Lemma 9. *If a vertex $\langle q, l \rangle$ is accepting, then it has rank $\alpha + 1$. Furthermore, if it has rank $\lambda + 1$, then it is accepting.*

Proof. The first part follows from the fact that an accepting vertex $\langle q, l \rangle$ is not inert: Therefore, $\langle q, l \rangle$ is a dead end in G_α , so $\langle q, l \rangle \notin G_{\alpha+1}$ and $\text{rank}_q^w(q, l) = \alpha + 1$.

For the second part, assume $\text{rank}_q^w(q, l) = \lambda + 1$, i.e., $\langle q, l \rangle \in G_\lambda \setminus G_{\lambda+1}$. Therefore, $\langle q, l \rangle$ is a dead end in G_λ . Since $G_\lambda \subseteq H_\lambda$, $\langle q, l \rangle$ is in H_λ as well. But H_λ has no dead ends, therefore $\langle q, l \rangle$ has at least one successor $\langle q', l + 1 \rangle$ in H_λ . But $\langle q, l \rangle$ is a dead end in G_λ , therefore any such successor $\langle q', l + 1 \rangle$ is not in G_λ . Therefore, $\langle q', l + 1 \rangle$ is inert in H_λ .

By contradiction, assume $\langle q, l \rangle$ that is not accepting. Since it has only inert successors $\langle q', l + 1 \rangle$ in H_λ , it is itself inert in H_λ . But $\langle q, l \rangle \in G_\lambda$, so $\langle q, l \rangle$ is *not* inert in H_λ . This is a contradiction, therefore $\langle q, l \rangle$ is accepting. \square

We say that a vertex $\langle q', l + 1 \rangle$ is a *maximal successor* of $\langle q, l \rangle$ if its rank is maximal amongst all successors of $\langle q, l \rangle$, and a sequence $\langle q_0, l \rangle \langle q_1, l + 1 \rangle \cdots \langle q_h, l + h \rangle$ is a *maximal path* if, for any $0 \leq k < h$, $\langle q_{k+1}, l + k \rangle$ is a maximal successor of $\langle q_k, l + k - 1 \rangle$.

We define a predecessor and a floor operation on ordinals. For an ordinal α , its predecessor $\alpha - 1$ is either α itself if α is a limit ordinal, or β if $\alpha = \beta + 1$ for some β ; its floor $\lfloor \alpha \rfloor := \sup_{\lambda < \alpha} \lambda$ is the largest limit ordinal strictly smaller than α . Notice that, for $0 < \alpha < \omega^\omega$, $\lfloor \alpha \rfloor < \alpha$.

Lemma 10. *Let vertex $\langle q, l \rangle$ have rank α . Then, a) every successor $\langle q', l + 1 \rangle$ has rank at most $\alpha - 1$, and b) there exists a maximal successor attaining rank $\alpha - 1$. As a direct consequence, c) every node $\langle q', l' \rangle$ reachable from $\langle q, l \rangle$ has a smaller rank $\alpha' \leq \alpha$.*

Proof. We split the proof in two cases, depending on whether α is a successor or limit ordinal. Let α be a successor ordinal $\beta + 1$. Then, $\langle q, l \rangle$ is a dead end in G_β , and thus it has no successor in G_β . Therefore, each successor $\langle q', l + 1 \rangle$ has rank $\leq \beta$. Moreover, we show that at least one successor has rank exactly equal to β . To this end, let $\beta^* \leq \beta$ be the maximum rank amongst $\langle q, l \rangle$'s successors. Notice that no successor $\langle q', l + 1 \rangle$ is in G_{β^*} . As $G_\beta \subseteq G_{\beta^*}$, it follows that $\langle q, l \rangle$ is a dead end in G_{β^*} . Therefore, $\langle q, l \rangle$ is not in G_{β^*+1} , which implies it has rank at most $\beta^* + 1 \leq \beta + 1$. But $\text{rank}_q^w(\langle q, l \rangle) = \beta + 1$ by assumption. Therefore, $\beta^* = \beta$, as required.

Otherwise, let α be a limit ordinal λ . Thus, $\langle q, l \rangle$ is inert in H_λ . Let $\langle q', l + 1 \rangle$ be a successor of $\langle q, l \rangle$. If $\langle q', l + 1 \rangle$ is not in H_λ , then, since $G_\lambda \subseteq H_\lambda$, $\langle q', l + 1 \rangle$ is not in G_λ either. Thus, $\langle q', l + 1 \rangle$ has rank $\leq \lambda$ in this case. Otherwise, let $\langle q', l + 1 \rangle$ be in H_λ . Since $\langle q, l \rangle$ is inert in H_λ , it follows that $\langle q', l + 1 \rangle$ is inert in H_λ as well. Therefore, $\langle q', l + 1 \rangle$ gets rank exactly equal to λ in this case. Finally, since H_λ does not contain dead ends, there exists at least one such inert successor in H_λ . \square

Lemma 11. *If a vertex $\langle q_0, l \rangle$ has a successor ordinal rank $\alpha + 1$, then there exists a maximal path $\langle q_0, l \rangle \langle q_1, l + 1 \rangle \cdots \langle q_h, l + h \rangle$ ending in $\langle q_h, l + h \rangle$ of rank $\lambda + 1$ with $\lfloor \alpha + 1 \rfloor \leq \lambda$.*

Proof. We proceed by ordinal induction. If α is a limit ordinal λ , the claim holds immediately: Take $h = 0$; clearly, $\lambda = \lfloor \lambda + 1 \rfloor$.

Otherwise, let α be a successor ordinal $\beta + 1$. That is, vertex $\langle q_0, l \rangle$ has rank $\alpha + 1 = (\beta + 1) + 1$. By Lemma 10 b), $\langle q_0, l \rangle$ has a maximal successor $\langle q_1, l + 1 \rangle$ of rank $\beta + 1 = \alpha$. By induction, there exists a maximal path $\langle q_1, l + 1 \rangle \cdots \langle q_h, l + h \rangle$ with $h > 0$, ending in $\langle q_h, l + h \rangle$ of rank $\lambda + 1$ with $\lfloor \beta + 1 \rfloor \leq \lambda$. But $\beta = \alpha - 1$, thus $\lfloor \beta + 1 \rfloor = \lfloor \alpha \rfloor \leq \lambda$. \square

Lemma 12. *If a vertex $\langle q_0, l \rangle$ has a nonzero limit ordinal rank λ , then there exists a path $\langle q_0, l \rangle \langle q_1, l + 1 \rangle \cdots \langle q_h, l + h \rangle$ with $h \geq 1$ ending in $\langle q_h, l + h \rangle$ of rank $\alpha + 1$ with $\lfloor \lambda \rfloor \leq \alpha$.*

Proof. Let $\langle q_0, l \rangle$ have rank $\lambda > 0$. By contradiction, assume $\langle q_0, l \rangle$ has no descendant $\langle q', l' \rangle$ of rank $\alpha + 1$ with $\lfloor \lambda \rfloor \leq \alpha$. That is, all descendants $\langle q', l' \rangle$ of successor ordinal rank $\alpha + 1$ have $\alpha < \lfloor \lambda \rfloor$, which is the same as $\alpha + 1 < \lfloor \lambda \rfloor$. By definition, $\langle q_0, l \rangle$ is inert in $H_\lambda \subseteq H_{\lfloor \lambda \rfloor}$. We show that $\langle q_0, l \rangle$ is inert in $H_{\lfloor \lambda \rfloor}$ as well. This is a contradiction, since λ is nonzero, therefore $\langle q_0, l \rangle$ would get rank $\lfloor \lambda \rfloor < \lambda$.

To this end, we show that any vertex reachable from $\langle q_0, l \rangle$ in $H_{\lfloor \lambda \rfloor}$ is non-accepting. For such a vertex $\langle q', l' \rangle$ to be accepting, by Lemma 9, it is necessary to have successor rank $\alpha + 1 < \lfloor \lambda \rfloor$. Clearly, $\langle q', l' \rangle \notin H_{\lfloor \lambda \rfloor}$. Therefore, $\langle q_0, l \rangle$ is inert in $H_{\lfloor \lambda \rfloor}$. \square

Lemma 13. *Let $w \in \Sigma^\omega$. If $\text{rank}_{q_0}^w(q_0, 0) \leq \text{rank}_{s_0}^w(s_0, 0)$, then $q_0 \sqsubseteq_w^{\text{de}} s_0$.*

Proof. Assume $\text{rank}_{q_0}^w(q_0, 0) \leq \text{rank}_{s_0}^w(s_0, 0)$. We show that Duplicator has a winning strategy in $G_w^{\text{de}}(q_0, s_0)$. For any round i , let $\langle q_i, s_i \rangle$ be the current configuration of the simulation game, and let the rank of Spoiler and Duplicator at round i be $\text{rank}_{q_0}^w(q_i, i)$ and $\text{rank}_{s_0}^w(s_i, i)$, respectively. Intuitively, Duplicator wins by ensuring both a *safety* and a *liveness* condition. The safety condition requires Duplicator to always preserve the ordering between ranks. I.e., at round i , $\text{rank}_{q_0}^w(q_i, i) \leq \text{rank}_{s_0}^w(s_i, i)$. The liveness condition enforces Duplicator to (eventually) visit an accepting state if Spoiler does so.

Duplicator plays in two modes, *normal mode* and *obligation mode*. In normal mode Duplicator only enforces the safety condition, while in obligation mode Duplicator needs to satisfy the liveness condition, while still preserving the safety condition.

In normal mode, we assume that Duplicator's rank is a limit ordinal, and, by Lemma 10, Duplicator can preserve the rank by always selecting maximal successors. We say that Duplicator *plays maximally* during normal mode. The game stays in normal mode as long as Spoiler is not accepting. Whenever $q_i \in F$ at round i , then Duplicator switches to obligation mode. Suppose that the current rank of Duplicator at round i is a limit ordinal λ . Since $q_i \in F$, by Lemma 9 Spoiler's rank is a successor ordinal $\alpha + 1 < \lambda$. W.l.o.g. we assume that Spoiler plays maximally during obligation mode. By Lemma 11, there exists a maximal path $\langle q_i, i \rangle \langle q_{i+1}, i+1 \rangle \cdots \langle q_j, j \rangle$ s.t. Spoiler's rank at round $j \geq i$ is $\lambda' + 1$. A further move by Spoiler extends the previous path to $\langle q_{j+1}, j+1 \rangle$. By Lemma 10 b), Spoiler's rank at round $j+1$ is now λ' , and by part c) of the same lemma, $\lambda' \leq \alpha + 1$. By part b), Duplicator can play a maximal path $\langle s_i, i \rangle \langle s_{i+1}, i+1 \rangle \cdots \langle s_{j+1}, j+1 \rangle$ s.t. Duplicator's rank at round $j+1$ is λ . Thus, $\lambda' < \lambda$, which implies $\lambda' \leq \lfloor \lambda \rfloor$. So, let $\langle q_{j+1}, s_{j+1} \rangle$ be the configuration at round $j+1$. By Lemma 12, Duplicator can play a path $\langle s_{j+1}, j+1 \rangle \langle s_{j+2}, j+2 \rangle \cdots \langle s_k, k \rangle$ with $k > j+1$ and s.t. Duplicator's rank at round k is $\alpha' + 1$ with $\lfloor \lambda \rfloor \leq \alpha'$. Therefore, $\lambda' \leq \alpha'$. By Lemma 11, Duplicator can extend the previous path with a maximal path $\langle s_k, k \rangle \langle s_{k+1}, k+1 \rangle \cdots \langle s_h, h \rangle$ s.t. Duplicator's rank at round $h > k$ is $\lambda'' + 1$ with $\lfloor \alpha' + 1 \rfloor \leq \lambda''$. By Lemma 9, $s_h \in F$, thus Duplicator has satisfied the pending obligation. At round $h+1$, Duplicator's rank is λ'' by Lemma 10 b), and the game can switch to normal mode. Notice that $\lambda' \leq \alpha' < \alpha' + 1$ implies $\lambda' \leq \lfloor \alpha + 1 \rfloor$. Therefore, $\lambda' \leq \lambda''$ and the safety condition is satisfied. \square

Lemma 14. *Let $w \in \Sigma^\omega$ and $k \geq 1$. If $q_0 \sqsubseteq_w^{k\text{-de}} s_0$, then $\text{rank}_{q_0}^w(q_0, 0) \leq \text{rank}_{s_0}^w(s_0, 0)$.*

Proof. We prove the contrapositive. Assume $\text{rank}_{q_0}^w(q_0, 0) \not\leq \text{rank}_{s_0}^w(s_0, 0)$. Since ordinals are linearly ordered, this means $\text{rank}_{q_0}^w(q_0, 0) > \text{rank}_{s_0}^w(s_0, 0)$. We have to show $q_0 \not\sqsubseteq_w^{k\text{-de}} s_0$, for arbitrary $k \geq 1$. Take n to be the size of the automaton. We actually prove that Duplicator does not win even with n pebbles, i.e., $q_0 \not\sqsubseteq_w^{n\text{-de}} s_0$.

For any round i , let $\langle q_i, s_i \rangle$ be the current configuration of the simulation game $G_w^{n\text{-de}}(q_0, s_0)$. (For simplicity, we omit the third component.) Notice that s_i identifies a subset of vertices at level i in $G_{s_0}^w$: $s_i \subseteq \{s \mid \langle s, i \rangle \in G_{s_0}^w\}$. We extend the notion of rank to sets of vertices by taking the maximal rank. That is, the rank of Duplicator at round i is $\sup_{s \in s_i} \text{rank}_{s_0}^w(s, i)$. As before, Spoiler's rank is just $\text{rank}_{q_0}^w(q_i, i)$.

We assume that, at round 0, every pebble has limit rank. If not, Spoiler can enforce such a situation by waiting a suitable number of rounds. (I.e., by playing maximally

according to Lemma 10.) So, let's Spoiler have limit rank λ and Duplicator have limit rank μ , with $\lambda > \mu$. We assume that Duplicator always plays maximally, unless she is forced to act differently. By Lemma 12, Spoiler can play a path $\langle q_0, 0 \rangle \langle q_1, 1 \rangle \cdots \langle q_i, i \rangle$ with $i > 0$, s.t. her rank at round i is $\alpha + 1$ and $\alpha \geq \lfloor \lambda \rfloor$. From $\lambda > \mu$ we have $\lfloor \lambda \rfloor \geq \mu$, which implies $\alpha \geq \mu$. By Lemma 11, Spoiler can extend the previous path with a maximal path $\langle q_i, i \rangle \langle q_{i+1}, i+1 \rangle \cdots \langle q_j, j \rangle$ with $j > i$, s.t. her rank at round j is $\lambda' + 1$ and $\lambda' \geq \lfloor \alpha + 1 \rfloor$. By Lemma 9, $q_j \in F$. From $\alpha + 1 > \alpha \geq \mu$ we have $\lfloor \alpha + 1 \rfloor \geq \mu$, which implies $\lambda' \geq \mu$. By performing a further maximal step, Spoiler reaches state $\langle q_{j+1}, j+1 \rangle$, thus attaining rank λ' . From now on, Spoiler plays maximally.

Since Duplicator was playing maximally, in the meanwhile she replied to Spoiler with a sequence $\langle s_0, 0 \rangle \langle s_1, 1 \rangle \cdots \langle s_{j+1}, j+1 \rangle$ s.t. she has rank μ at round $j+1$.

Now, let $\langle q_{j+1}, s_{j+1} \rangle$ be the current configuration, and remember that Duplicator has a pending obligation. That is, Duplicator has to ensure that at some future round k all pebbles are good since round $j+1$. Let s_k be the position of pebbles at round k . This implies that every state in s_k has an accepting predecessor since round $j+1$. By Lemma 9, accepting pebbles receive successor ranks, and, since ranks are nonincreasing along paths in $G_{s_0}^w$ (by Lemma 10), it follows that every pebble in s_k has rank $< \mu$. That is, Duplicator's rank at round k is $< \mu$. Since Duplicator has now satisfied the pending obligation, she will again play maximally, from round k on. By Lemma 10, all pebbles eventually stabilize to a limit rank. Since there is a finite number of pebbles, it follows that at some round $h \geq k$ Duplicator's rank is $\mu' < \mu$. Let s_h be the position of Duplicator's pebbles at round h .

In the meanwhile Spoiler replied with a maximal path $\langle q_{j+1}, j+1 \rangle \cdots \langle q_h, h \rangle$, preserving rank $\lambda' \geq \mu > \mu'$ until round h . Therefore, $\lambda' > \mu'$ and the situation at round h is identical to the initial situation at round 0.

Since ordinals are well-founded, Spoiler can iterate the whole procedure and after a finite number of repetitions Duplicator hits the trap rank ω . At that point, Spoiler would have a limit rank $\lambda' > \omega$, so she will just force one more obligation, which would remain unmet (vertices of rank ω have no accepting successor). Thus, Spoiler wins. \square

Theorem 3. For any NBA \mathcal{Q} , $k \geq 1$ and states $q, s \in \mathcal{Q}$, $q \sqsubseteq_{fx}^{k-de} s$ iff $q \sqsubseteq_{fx}^{de} s$.

Proof. By combining the previous two lemmas, we get

$$q \sqsubseteq_{fx}^{de} s \implies q \sqsubseteq_{fx}^{k-de} s \implies (\text{rank}_q^w(q, 0) \leq \text{rank}_s^w(s, 0)) \implies q \sqsubseteq_{fx}^{de} s,$$

where the first implication holds by the definition of \sqsubseteq_{fx}^{k-de} , and the last two by Lemmas 14 and 13, respectively.

B Proofs for Section 4

Lemma 3. Let $w \in \Sigma^\omega$ and π_0, π_1, \dots as in Definition 1. If $\Pi = \pi_0, \pi_1, \dots$ is coherent, then any infinite subsequence $\Pi' = \pi_{f(0)}, \pi_{f(1)}, \dots$ thereof is coherent.

Proof. Let $\Pi := \pi_0, \pi_1, \dots$ be an infinite coherent sequence, and let $\Pi' := \pi_{f(0)}, \pi_{f(1)}, \dots$ be any infinite subsequence thereof, for some $f : \mathbb{N} \mapsto \mathbb{N}$ with $f(0) < f(1) < \dots$. We

have to show

$$\forall i' \cdot \exists j' \cdot \exists h' \cdot \forall k' \geq h' \cdot j' < |\pi_{f(k')}| \wedge \text{cnt-final}(\pi_{f(k')}, j') \geq i'.$$

Let $i' \in \mathbb{N}$. By taking $i := i'$, by the coherence of Π , there exists j, h s.t

$$(*) \forall k \geq h \cdot j < |\pi_k| \wedge \text{cnt-final}(\pi_k, j) \geq i'.$$

Let h' be the minimal m s.t. $f(m) \geq h$. For any $k' \geq h'$, we have $f(k') \geq f(h') \geq h$. Thus, by letting $k := f(k')$ in $(*)$, we obtain $j < |\pi_{f(k')}| \wedge \text{cnt-final}(\pi_{f(k')}, j) \geq i'$. Take $j' := j$. Since $k' \geq h'$ was arbitrary, we have proved that Π' is coherent. \square

Lemma 4. For $w \in \Sigma^\omega$, let $\Pi = \pi_0, \pi_1, \dots$ be a coherent sequence of paths over (prefixes of) w . Then, there exists a fair path ρ over w . Moreover, if all π_i 's are initial, then ρ is initial.

Proof. Let $\Pi := \pi_0, \pi_1, \dots$ be a coherent sequence. We prove by induction the following claim: For $l \in \mathbb{N}$, $R(l)$ holds iff there exists a finite sequence of finite paths $\rho_0 <_{\text{prf}} \rho_1 <_{\text{prf}} \dots <_{\text{prf}} \rho_l$, with ρ_l of length m_l , and an infinite subsequence $\Pi_l := \pi_{f_l(0)}, \pi_{f_l(1)}, \dots$ of Π with $f_l(0) < f_l(1) < \dots$, such that

$$(a) \text{cnt-final}(\rho_l, m_l) \geq l \quad (b) \Pi_l \text{ is coherent} \quad (c) \forall k \cdot \rho_l \leq_{\text{prf}} \pi_{f_l(k)}. \quad (3)$$

For the base case $l = 0$, take $\rho_0 := \varepsilon$ of length $m_0 := 0$, and $f_0(i) = i$ for any i . Then, $\Pi_0 = \Pi$ and $R(0)$ holds.

For the inductive step, assume $R(l-1)$ holds. That is, there exist $\rho_0 <_{\text{prf}} \rho_1 <_{\text{prf}} \dots <_{\text{prf}} \rho_{l-1}$, with ρ_{l-1} of length m_{l-1} , and $\Pi_{l-1} = \pi_{f_{l-1}(0)}, \pi_{f_{l-1}(1)}, \dots$ with $\rho_{l-1} \leq_{\text{prf}} \pi_{f_{l-1}(k)}$ for any k . Since Π_{l-1} is coherent, by taking $i := l$, there exist j and h s.t., for any π in the sequence $\pi_{f_{l-1}(h)}, \pi_{f_{l-1}(h+1)}, \dots$, π has length at least j and $\text{cnt-final}(\pi, j) \geq l$. Since the various π 's are branches in a finitely-branching tree, it follows that at any fixed depth d there are only finitely many different branches of length d . Therefore, there exists a least one such finite branch which is shared by infinitely many π 's. For $d = j$, we get that there exists a finite path ρ' of length j s.t. $\text{cnt-final}(\rho', j) \geq l$ and $\rho' \leq_{\text{prf}} \pi$ for infinitely many such π 's. Let $\Pi_l := \pi_{g(f_{l-1}(h))}, \pi_{g(f_{l-1}(h+1))}, \dots$ be this infinite subsequence. We assume w.l.o.g. that $m_{l-1} < j$, and, consequently, $\rho_l <_{\text{prf}} \rho'$. Take $f_l := g \circ f_{l-1}$, $\rho_l := \rho'$ and $m_l := j$. Then, (a) and (c) are satisfied by construction, while (b) follows by Lemma 3. This proves $R(l)$, concluding the inductive step.

Therefore, one can build the infinite sequence of finite paths $\varepsilon = \rho_0 <_{\text{prf}} \rho_1 <_{\text{prf}} \dots$ such that, for any l , ρ_l visits at least l final states. Take ρ to be the limit of the ρ_l 's. Finally, since $\rho_1 <_{\text{prf}} \pi_{f_1(0)}$ by property (c), it follows that if all π_i 's are initial, then so is $\pi_{f_1(0)}$, and thus ρ . \square

Theorem 4. Let R be a jumping-safe preorder. Then, R is good for quotienting.

Proof. Assume R is jumping-safe and let \approx_R be the equivalence induced by R . We have to show $\mathcal{L}^\omega(\mathcal{Q}) = \mathcal{L}^\omega(\mathcal{Q}_{\approx_R})$. The direction $\mathcal{L}^\omega(\mathcal{Q}) \subseteq \mathcal{L}^\omega(\mathcal{Q}_{\approx_R})$ holds by Lemma 1.

For the other direction, assume $w \in \mathcal{L}^\omega(\mathcal{Q}_{\approx_R})$, with $w = a_0 a_1 \dots \in \Sigma^\omega$. Let $\pi_{\approx_R} = [q_0] \xrightarrow{a_0} [q_1] \xrightarrow{a_1} [q_2] \dots$ be an accepting run over w in \mathcal{Q}_{\approx_R} . By the definition of quotient, for any i , there exist states $q_i, q_i^F, \hat{q}_i \in Q$ s.t. $q_i R q_i^F R \hat{q}_i$ and $\hat{q}_i \xrightarrow{a_i} q_{i+1}$. That is, π_{\approx_R} induces a jumping path π as in Equation 2. Moreover, q_i^F can be taken in F if $[q_i]$ is accepting. Since $[q_0]$ is initial, we assume w.l.o.g. that $q_0 \in I$. Since R is jumping-safe and π is both initial and fair, there exists a coherent sequence of initial paths π_0, π_1, \dots over prefixes of w . By Lemma 4, there exist an (non-jumping) accepting run over w in \mathcal{Q} . Therefore, $w \in \mathcal{L}^\omega(\mathcal{Q})$. \square

C Proofs for Section 5

Lemma 5. *For a preorder R , $R \subseteq R \circ \tau_0(R) \subseteq \tau_0(R)$.*

Proof. Directly from Lemmas 15 and 16 below.

Lemma 15. *For any reflexive R , $R \subseteq \tau_0(R)$.*

Proof. Let $T := \tau_0(R)$, and assume $s R q$. We have to show $s T q$. Let's Spoiler select a and q' s.t. $q \xrightarrow{a} q'$. Since $s R q$ by assumption, Duplicator can directly take $\hat{s} := q$. Trivially $q \in F \implies \hat{s} \in F$, as required by the winning condition. \square

Lemma 16. *For any transitive R , $R \circ \tau_0(R) \subseteq \tau_0(R)$.*

Proof. Let $T := \tau_0(R)$, and assume $\bar{s} R s T q$. We have to show $\bar{s} T q$. Let's Spoiler select a and q' s.t. $q \xrightarrow{a} q'$. Since $s T q$ by assumption, Duplicator can select \hat{s} s.t. $s R \hat{s}$ and $\hat{s} \xrightarrow{a} s'$, for some s' . Then, by transitivity, $\bar{s} R \hat{s}$. As $q \in F \implies \hat{s} \in F$ (by $s T q$), we conclude that Duplicator wins from \bar{s} as well, thus $\bar{s} T q$. \square

Theorem 5. *Let R a F -respecting preorder, and let $T \subseteq \tau_0(R)$ be an appealing, improving fragment of $\tau_0(R)$. If R is jumping-safe, then T is jumping-safe.*

Proof. Assume that R is jumping-safe and F -respecting, and let T be an appealing, improving fragment of $\tau_0(R)$. That is, T is a self-respecting and transitive fragment of $\tau_0(R)$, with $R \subseteq T$. We have to show that T is jumping-safe. To this end, let $w = a_0 a_1 \dots \in \Sigma^\omega$, and let the following be an initial T -jumping path

$$\pi = q_0 T q_0^F T \hat{q}_0 \xrightarrow{a_0} q_1 T q_1^F T \hat{q}_1 \xrightarrow{a_1} q_2 \dots, \quad q_0 \in I.$$

First, we show by induction the following claim: For any $i \geq 0$, there exists a finite initial path

$$\rho_i = r_0 R \hat{r}_0 \xrightarrow{a_0} r_1 R \hat{r}_1 \xrightarrow{a_1} \dots r_i, \quad r_0 \in I,$$

s.t. $r_i T q_i$, and, for any $0 \leq k < i$, $q_k^F \in F \implies \hat{r}_k \in F$.

For $i = 0$, just take $r_0 := q_0$. For $i \geq 0$, assume $\rho_i = r_0 R \hat{r}_0 \xrightarrow{a_0} r_1 \dots r_i$ has already been built. Since $q_i^F T \hat{q}_i \xrightarrow{a_i} q_{i+1}$, by the definition of T there exists $\hat{q}_i^F \xrightarrow{a_i} q'$ for some \hat{q}_i^F and q' with $q_i^F R \hat{q}_i^F$ and $q' T q_{i+1}$. But $q_i T q_i^F$ and, by induction hypothesis, $r_i T q_i$. Since T is transitive, we get $r_i T q_i^F$, so there exists

$\hat{r}_i \xrightarrow{a_i} r_{i+1}$ with $r_i R \hat{r}_i$ and $r_{i+1} T q'$. Again by transitivity, we get $r_{i+1} T q_{i+1}$. Moreover, if $q_i^F \in F$, then since R respects final states, we have $\hat{q}_i^F \in F$, and, by the definition of T , we finally derive $\hat{r}_i \in F$. Thus, we have just built $\rho_{i+1} = r_0 R \hat{r}_0 \xrightarrow{a_0} r_1 \cdots r_i R \hat{r}_i \xrightarrow{a_i} r_{i+1}$. This concludes the inductive step, and the claim is proved.

From the claim above, let ρ be the infinite initial R -jumping sequence resulting by taking limit of the ρ_i 's. Since R is jumping-safe, there exists an infinite sequence of initial finite paths π_0, π_1, \dots s.t. $\text{last}(\pi_i) R r_i$. By assumption $R \subseteq T$, so $\text{last}(\pi_i) T r_i$ holds as well. By $r_i T q_i$ and transitivity, we obtain $\text{last}(\pi_i) T q_i$. Therefore, the same sequence π_0, π_1, \dots can be taken as a witness for T being jumping-safe.

Finally, assume that π is fair, i.e., $q_i^F \in F$ for infinitely many i 's. By the claim above, $\hat{r}_i \in F$ for infinitely many i 's, therefore ρ is fair as well. Since R is jumping-safe (by taking $r_i^F := \hat{r}_i$, R being reflexive), we finally infer that π_0, π_1, \dots is coherent, which concludes the proof. \square

Lemma 6. *For any reflexive R , let $T \subseteq \tau_0(R)$ be any appealing fragment of $\tau_0(R)$. Then, $\tau_0(T) \subseteq \tau_0(R)$. That is, at the second iteration τ_0 does not introduce any new fragment which could not be found before.*

Proof. Let R be reflexive. Let T be an appealing (= transitive and self-respecting) fragment of $V_0 := \tau_0(R)$, and let $V_1 := \tau_0(T)$. We have to show $V_1 \subseteq V_0$. To this end, let $s V_1 q$ and let Spoiler choose a transition $q \xrightarrow{a} q'$. By the definition of V_1 , there exist $s T \bar{s}$ and \bar{s}' with $\bar{s} \xrightarrow{a} \bar{s}'$ and $\bar{s}' V_1 q'$. By the definition of T , there exist $s R \hat{s}$ and $s' T \bar{s}'$ (since T is self-respecting). T being transitive, from $s' V_1 \bar{s}' T q'$ and from Lemma 16, we get $s' V_1 q'$. Thus, we let Duplicator choose \hat{s} and s' above, as required by the definition of V_0 . Duplicator is winning as $q \in F$ implies $\bar{s} \in F$, and the latter implies $\hat{s} \in F$, the first implication holding by the definition of V_1 , and the second by T . Therefore, $s V_0 q$. \square

Lemma 17. *For any relation R , $\tau_1(R)$ is transitive.*

Proof. Let $T := \tau_1(R)$, and let $s T r T p$. We have to show $s T p$. Let Spoiler choose a and \hat{p} and p' s.t. $p R \hat{p}$ and $\hat{p} \xrightarrow{a} p'$. We have to show 1) that Duplicator can choose \hat{s} and s' s.t. $s R \hat{s}$ and $\hat{s} \xrightarrow{a} s'$, and 2) $\hat{p} \in F \implies \hat{s} \in F$. For 1), from $r T p$ it follows that there exist \hat{r} and r' s.t. $r R \hat{r}$ and $\hat{r} \xrightarrow{a} r'$. Then, from $s T r$ one can directly find the required \hat{s} and s' . For 2), assume $\hat{p} \in F$. From $r T p$ it follows that the \hat{r} found above is in F as well. Finally, $\hat{s} \in F$ follows from $s T r$ in a similar way. \square

Lemma 18. *For any transitive R , $R \subseteq \tau_1(R)$.*

Proof. Let $T := \tau_1(R)$, and assume $s R q$. We have to show $s T q$. Let's Spoiler select a and \hat{q} and q' s.t. $q R \hat{q}$ and $\hat{q} \xrightarrow{a} q'$. Since $s R q$ by assumption, and from R being transitive, we have $s R \hat{q}$. Thus Duplicator can directly take $\hat{s} := \hat{q}$. Finally, trivially $\hat{q} \in F \implies \hat{s} \in F$, as required by the winning condition. \square

Lemma 7. For any R , let $T \subseteq \tau_0(R)$ be any appealing fragment of $\tau_0(R)$. If $R \subseteq T$ (i.e., R is improving), then $T \subseteq \tau_1(R)$.

Proof. Let R, T as in the statement of the lemma, and let $V := \tau_1(R)$. We have to show $T \subseteq V$. Let $q T p$, and let Spoiler choose \hat{p} and p' with $p R \hat{p}$ and $\hat{p} \xrightarrow{a} p'$, as required by the definition of V . Then, as $R \subseteq T$ by assumption, and T being transitive, we have $q T \hat{p}$. Therefore, by the definition of T , Duplicator can choose \hat{q} and q' with $q R \hat{q}$ and $\hat{q} \xrightarrow{a} q'$. Since T is self-respecting, we have $p' T q'$. Finally, $\hat{p} \in F \implies \hat{s} \in F$ by the definition of T . Therefore, Duplicator is winning, and $q V p$. \square




Fig. 7. Quotienting w.r.t. appealing fragments of τ_0^{de} is incorrect, already for unary automata. We have $q_3 \sqsubseteq_{\text{bw}}^{\text{di}} q_2$ and $q_4 \sqsubseteq_{\text{bw}}^{\text{di}} q_2$, and the relation $T := \{(q_i, q_i) \mid 0 \leq i \leq 6\} \cup \{(q_i, q_6) \mid 0 \leq i \leq 6\} \cup \{(q_i, q_j) \mid 2 \leq i, j \leq 4\} \cup \{(q_i, q_5) \mid 2 \leq i \leq 4\}$ is an appealing fragment of $\tau_0^{\text{de}}(\sqsubseteq_{\text{bw}}^{\text{di}})$. (In particular, $q_3 \tau_0^{\text{de}}(\sqsubseteq_{\text{bw}}^{\text{di}}) q_4$ since q_3 can “jump” to q_2 .) The equivalence induced by T identifies the states q_2, q_3, q_4 , but this is incorrect as the resulting automaton would accept the spurious word a^ω .

Lemma 8. For any R , $\tau_1^{\text{de}}(R)$ is transitive.

Proof. A complete and formal proof of transitivity requires the machinery of logbooks and composition of (winning) strategies, which is a standard tool for delayed simulation (for more details see, e.g., [7]). Here, we highlight the ingredients pertinent to τ_1^{de} .

Let $T := \tau_1^{\text{de}}(R)$, and let $r T q T p$. We have to show $r T p$. Let G_0 be the game between r and q , let G_1 be the game between q and p , and let G be the outer game between r and p .

The idea is that Duplicator plays G and at the same time updates G_0, G_1 accordingly. At round i , if the G -configuration is $\langle r_i, p_i \rangle$, then there exists q_i s.t. the G_0 configuration is $\langle r_i, q_i \rangle$ and the G_1 configuration is $\langle q_i, p_i \rangle$.

Let Spoiler choose \hat{p} and transition $\hat{p} \xrightarrow{a_i} p_{i+1}$, with $p_i R \hat{p}$. Since G_1 -Duplicator is winning, there exist \hat{q} and transition $\hat{q} \xrightarrow{a_i} q_{i+1}$, with $q_i R \hat{q}$. Similarly, G_0 -since Duplicator is winning, there exist \hat{r} and transition $\hat{r} \xrightarrow{a_i} r_{i+1}$, with $r_i R \hat{r}$. Thus, Duplicator can proceed in G by taking the last transition above. The configurations are updated as follows: The game G_0 goes to $\langle r_{i+1}, q_{i+1} \rangle$, G_1 goes to $\langle q_{i+1}, p_{i+1} \rangle$ and G goes to $\langle r_{i+1}, p_{i+1} \rangle$.

We now argue that the strategy above is winning. W.l.o.g. we assume that the games G_0, G_1 are updated according to a fixed winning strategy. We show that Duplicator is

winning in G . Assume $\hat{p}_i \in F$. Since G_1 -Duplicator is playing according to a winning strategy, there exists $k \geq i$ s.t. $\hat{q}_k \in F$. Similarly, as G_0 -Duplicator is playing according to a winning strategy, there exists $j \geq k \geq i$ s.t. $\hat{r}_j \in F$. Thus, take $j \geq i$ s.t. $\hat{r}_j \in F$, as required. \square

Lemma 19. *For any transitive R , $R \subseteq \tau_1^{de}(R)$.*

Proof. Immediate from $R \subseteq \tau_1(R)$ by Lemma 18, and $\tau_1(R) \subseteq \tau_1^{de}(R)$ by definition. \square

Theorem 6. *If R is a jumping-safe F -respecting preorder, then $\tau_1^{de}(R)$ is jumping-safe.*

Proof. Assume that R is a jumping-safe, F -respecting preorder, and let $T := \tau_1^{de}(R)$. We have to show that T is jumping-safe. During the proof we refer to Figure 8, hereafter called “the diagram”. Let $w = a_0 a_1 \dots \in \Sigma^\omega$, and let π be an initial T -jumping path

$$\pi = q_0 \ T \ q_0^F \ T \ \hat{q}_0 \xrightarrow{a_0} q_1 \ T \ q_1^F \ T \ \hat{q}_1 \xrightarrow{a_1} q_2 \dots, \quad q_0 \in I.$$

See the blue path in the diagram. We inductively show how to build the rest of the diagram, and then we use this construction for showing that T is jumping-safe.

Formally, we inductively build a sequence $\rho_0, \rho_1, \dots, \rho_i$ such that, for any $k \leq i$, ρ_k is a T -ordered $k + 4$ -tuple of states representing the k -th layer of the diagram,

$$\rho_k = s_k^0 \ T \ s_k^1 \ T \ \dots \ T \ s_k^{k-1} \ T \ s_k^k \ T \ q_k \ T \ q_k^F \ T \ \hat{q}_k.$$

Two successive layers are in relations with transitions as follows (cf. the diagram):

$$\forall (1 \leq h \leq k) \cdot s_k^h \xrightarrow{a_k} s_{k+1}^h, \quad q_k^F \xrightarrow{a_k} s_{k+1}^{k+1}, \quad \hat{q}_k \xrightarrow{a_k} q_{k+1},$$

where the dashed arrow $x \xrightarrow{a} y$ represents an R -jumping transition via some suitable proxy. That is, $x \xrightarrow{a} y$ iff there exists a proxy \hat{x} s.t. $x \ R \ \hat{x}$ and $\hat{x} \xrightarrow{a} y$.

For $i = 0$, just take $s_0^0 := q_0$. Then, the invariant is clearly satisfied, as $q_0 \ T \ q_0^F$ by assumption and $q_0 \ T \ q_0$ by T being reflexive.

For $i \geq 0$, assume $\rho_0, \rho_1, \dots, \rho_i$ has already been built. By induction hypothesis, ρ_i is the following T -ordered tuple:

$$\rho_i = s_i^0 \ T \ s_i^1 \ T \ \dots \ T \ s_i^{i-1} \ T \ s_i^i \ T \ q_i \ T \ q_i^F \ T \ \hat{q}_i,$$

The next layer ρ_{i+1} ,

$$\rho_{i+1} = s_{i+1}^0 \ T \ s_{i+1}^1 \ T \ \dots \ T \ s_{i+1}^{i-1} \ T \ s_{i+1}^i \ T \ s_{i+1}^{i+1} \ T \ q_{i+1} \ T \ q_{i+1}^F \ T \ \hat{q}_{i+1},$$

is obtained as follows. The last three components $q_{i+1}, q_{i+1}^F, \hat{q}_{i+1}$ are fixed by the T -jumping path π . The rest is determined next. Since $\hat{q}_i \xrightarrow{a_i} q_{i+1}$, we propagate the transition down the chain, by using the definition of T —as indicated by the zigzag arrows in the diagram. As $q_i^F \ T \ \hat{q}_i$, there exists an R -jumping transition $q_i^F \xrightarrow{a_i} q' \ T \ q_{i+1}$. Take $s_{i+1}^{i+1} := q'$. Similarly, from $s_i^i \ T \ q_i \ T \ q_i^F$ there exists $s_i^i \xrightarrow{a_i} q'' \ T \ s_{i+1}^{i+1}$. Take $s_{i+1}^i := q''$. Clearly, one can build all the remaining states down to s_{i+1}^0 in the same way, thus completing layer $i + 1$ in the diagram. This concludes the inductive step in the definition of ρ_{i+1} .

Remark. We assume that each time a new T -game starts from configuration $\langle q_i^F, q_i \rangle$, Duplicator fixes a winning strategy, and always plays accordingly.

We now prove that final states are “propagated” in the diagram right-to-left, top-to-bottom: Formally, we show that, for any $i \geq 0$, if $q_i^F \in F$, then there exists $j \geq i$ s.t. $\hat{s}_j^0 \in F$, where \hat{s}_j^0 is the proxy witnessing $s_j^0 \xrightarrow{a_j} s_{j+1}^0$. Assume $q_i^F \in F$. Then, since R is F -respecting, $\hat{q}_i^F \in F$, where \hat{q}_i^F is the proxy witnessing $q_i^F \xrightarrow{a_i} s_{i+1}^{i+1}$. Since $s_i^i T q_i^F$, by the definition of τ_1^{de} and by the above remark, there exists $j_0 \geq i$ s.t. $\hat{s}_{j_0}^i \in F$, where $\hat{s}_{j_0}^i$ is the proxy witnessing $s_{j_0}^i \xrightarrow{a_{j_0}} s_{j_0+1}^i$. But $s_{j_0}^{i-1} T s_{j_0}^i$, therefore there exists $j_1 \geq j_0$ s.t. $\hat{s}_{j_1}^{i-1} \in F$, and so on ... until we reach index $j_i \geq j_{i-1}$, for which $\hat{s}_{j_i}^0 \in F$. Thus, take $j := j_i$.

We are finally ready to prove that T is jumping-safe. Notice that the leftmost path in the diagram represents an initial R -jumping path π' ,

$$\pi' = s_0^0 R \hat{s}_0^0 \xrightarrow{a_0} s_1^0 R \hat{s}_1^0 \xrightarrow{a_1} \dots, s_0^0 = q_0 \in I.$$

Since R is jumping-safe, there exists an infinite sequence of initial finite paths π_0, π_1, \dots s.t. $\text{last}(\pi_i) R s_i^0$. Since R is transitive, $R \subseteq T$ by Lemma 19. Therefore, $\text{last}(\pi_i) T s_i^0$. By $s_i^0 T q_i$ and transitivity, we obtain $\text{last}(\pi_i) T q_i$. Therefore, the same sequence π_0, π_1, \dots can be taken as a witness for T being jumping-safe. Finally, since π is fair, i.e., $q_i^F \in F$ for infinitely many i 's, then π' is fair, as final states are “propagated” (shown above). Since R is jumping-safe, we conclude that π_0, π_1, \dots is coherent. \square




Fig. 8. Construction for the proof of Theorem 6.

By using similar techniques, it is possible to show that repeated application of τ_1^{de} does not give coarser relations. This is analogous of what proved in Lemma 6 for τ_0 . The proof of this fact is omitted.

Lemma 20. For any preorder R , $\tau_1^{\text{de}}(\tau_1^{\text{de}}(R)) \subseteq \tau_1^{\text{de}}(R)$.

D Computing $\tau_1^{\text{de}}(R)$

In this section we give an algorithm for computing $\tau_1^{\text{de}}(R)$ from Section 5.1, obtained as an extension of the classical algorithm for computing delayed simulation [6]. We assume that the relation R has already been computed. We build a game graph where Duplicator has a Büchi winning objective.

We enrich configurations from the basic semantic game for $\tau_1^{\text{de}}(R)$ with an obligation bit recording whether Duplicator has to visit an accepting state. Formally, Spoiler's positions are of the form $\langle s, q, b \rangle$, with $q, s \in Q$ and $b \in \{0, 1\}$, and Duplicator's positions are of the form $\langle s, q, \hat{b}, a, q' \rangle$, with $q, s, q' \in Q$, $a \in \Sigma$ and $\hat{b} \in \{0, 1\}$. Spoiler can pick a move $(\langle s, q, b \rangle, \langle s, q, \hat{b}, a, q' \rangle) \in \Gamma'_0$ if there exists $\hat{q} \in Q$ s.t. $qR\hat{q} \xrightarrow{a} q'$, and $\hat{b} = 1$ if $\hat{q} \in F$ and b otherwise. Similarly, Duplicator can pick a move $(\langle s, q, \hat{b}, a, q' \rangle, \langle s', q', b' \rangle) \in \Gamma'_1$ if there exists $\hat{s} \in Q$ s.t. $sR\hat{s} \xrightarrow{a} s'$, and $b' = 0$ if $\hat{s} \in F$ and \hat{b} otherwise. The objective for Duplicator is to ensure that the winning bit is 0 infinitely often, that is, every obligation to visit an accepting state is eventually met. Formally, the winning condition is

$$W' = \{ \langle s_0, q_0, b_0 \rangle \langle s_1, q_1, b_1 \rangle \cdots \mid \forall i \geq 0 \cdot \exists j \geq i \cdot b_j = 0 \} .$$

Let CPre be a controlled predecessor operator for Duplicator, defined as

$$\text{CPre}(X) = \{ x \mid \forall (x, y) \in \Gamma'_0 \cdot \exists (y, z) \in \Gamma'_1 \cdot z \in X \} .$$

That is, $x = \langle s, q, b \rangle \in \text{CPre}(X)$ if Duplicator can force the game in X in one step from configuration x . Then, the winning region for Duplicator can be computed by evaluating the following fixpoint:

$$V = \nu X \cdot \mu Y \cdot [b = 0] \cap \text{CPre}(X) \cup \text{CPre}(Y) ,$$

where with $[b = 0]$ we have indicated the set of configurations with no obligation pending, i.e., $[b = 0] = \{ \langle q, s, b \rangle \mid q, s \in Q, b = 0 \}$. Finally, $s \tau_1^{\text{de}}(R) q$ holds iff $\langle s, q, 0 \rangle \in V$.

E Proof of Theorem 1

First, we define yet another refinement transformer, called *fixed-word delayed transformer* $\tau_1^{\text{fx-de}}$, which is the same as τ_1^{de} , with the only difference that Spoiler has to reveal the whole input word $w = a_0 a_1 \cdots$ in advance. Notice that $\tau_1^{\text{fx-de}}$, though not efficiently computable in general, has properties very similar to τ_1^{de} . In particular, the proof of Theorem 6 works as it is for the lemma below.

Lemma 21. If R is a jumping-safe F -respecting preorder, then, $\tau_1^{\text{fx-de}}(R)$ is jumping-safe.

Theorem 1. $\sqsubseteq_{\text{fx}}^{\text{de}}$ is good for quotienting.

Proof. Directly from Lemma 21, since $\sqsubseteq_{\text{fx}}^{\text{de}}$ is (the transpose of) $\tau_1^{\text{fx-de}}$ applied to the identity relation.